Signature-aided tracking using HRR profiles

Kevin J. Sullivan[*a], Matthew B. Ressler[b], Robert L. Williams[c]

[a]Toyon Research Corporation; [b]SRI, International; [c]Air Force Research Laboratory

ABSTRACT

This paper describes a methodology for using high-range-resolution ground-moving-target-indicator (HRRGMTI) profiles to aid the association process in an automated tracker. The tracker uses a Kalman filter to estimate the state of targets. These state estimates are used in addition to HRRGMTI measurements to perform data association. The HRRGMTI profiles used to develop and test the system are simulated using the MSTAR dataset of synthetic aperture radar (SAR) imagery. The performance of the system was estimated using a computer simulation that modeled the radars, ground vehicles, and creation of HRRGMTI profiles. Performance relative to a tracker that uses only kinematic information to perform data association is computed and presented.

Keywords: data association, tracking, automatic target recognition, ATR, HRRGMTI, range profiles.

1. INTRODUCTION

In most scenarios of interest to U.S. commanders, it will be critically beneficial to know the current location of certain types of ground vehicles for extended periods of time. For instance, knowing the location of tactical ballistic missile (TBM) launchers at all times allows commanders to plan strikes at the most appropriate time. Additionally, maintaining track of mobile surface-to-air missiles provides current knowledge of air defense locations. In low-intensity conflicts or operations-other-than-war, it would be useful to be able to track vehicles known to be carrying war criminals, weapons of mass destruction, or terrorists, so that they can be apprehended or attacked at the most feasible time. In larger-scale conflicts, U.S. commanders will want to know the locations of certain types of high-value military units such as armor, engineering units, or mobile command posts so that the intentions of enemy commanders can be inferred and so that the units can be attacked at a time that is most beneficial to U.S. commanders.

Synthetic aperture radars (SAR) and ground moving target indicator (GMTI) radars, can be used to detect vehicles while they are stationary and while they are moving, respectively. While development of advanced exploitation techniques for SAR and GMTI data continues, less work is being done to better fuse information from the sensors and task them in a cooperative fashion so that the location and identity of ground vehicles can be maintained. The use of SAR and GMTI to track vehicles over extended move-stop-move cycles has been called continuous tracking or the maintenance of continuous identification. We will use the term Continuous Identification (or Continuous ID) to stress the importance of being able to maintain an identification of the vehicle over an extended period of time.

Toyon Research Corporation has been involved in the development of the Continuous ID concept for many years.[1,2] We have been active in the development of algorithms in several key technology areas. These developments have come a long way toward a solution to the continuous ID problem, but a complete solution does not yet exist. The work described in this paper was supported by a phase I SBIR sponsored by AFRL/SN. In this effort, we have focused on the key technology of using high-range-resolution (HRR) GMTI signatures to aid the association of GMTI detections when forming vehicle tracks. These ideas are presented in section 2.1 through 2.3 of this paper. In order to evaluate the potential effectiveness of these ideas, we have developed a testbed that models the operation of the sensors, ground vehicles, and algorithms involved in Continuous ID. The testbed is described in section 2.4. We have exercised the testbed on an example problem and present the results of this analysis in section 2.5.

2. TECHNICAL DISCUSSION

2.1 Signature-aided-tracking (SAT) approach

There are two basic approaches that could be employed to perform SAT. The first technique relies on the ability to be able to classify target vehicles using stored templates that are created prior to the execution of a mission. Each new sensor

---

[*] ksullivan@toyon.com; phone 1 805 968-6787; fax 1 805 685-8089; http://www.toyon.com; Toyon Research Corporation, 75 Aero Camino, Suite A, Goleta, CA 93117-3139.

1. REPORT DATE **2006**	2. REPORT TYPE		3. DATES COVERED **00-00-2006 to 00-00-2006**
4. TITLE AND SUBTITLE **Signature-aided tracking using HRR profiles**			5a. CONTRACT NUMBER
			5b. GRANT NUMBER
			5c. PROGRAM ELEMENT NUMBER
6. AUTHOR(S)			5d. PROJECT NUMBER
			5e. TASK NUMBER
			5f. WORK UNIT NUMBER
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) **Toyon Research Corporation,75 Aero Camino Suite A,Goleta,CA,93117**			8. PERFORMING ORGANIZATION REPORT NUMBER
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)			10. SPONSOR/MONITOR'S ACRONYM(S)
			11. SPONSOR/MONITOR'S REPORT NUMBER(S)
12. DISTRIBUTION/AVAILABILITY STATEMENT **Approved for public release; distribution unlimited**			
13. SUPPLEMENTARY NOTES **The original document contains color images.**			
14. ABSTRACT **see report**			
15. SUBJECT TERMS			

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES **11**	19a. NAME OF RESPONSIBLE PERSON
a. REPORT **unclassified**	b. ABSTRACT **unclassified**	c. THIS PAGE **unclassified**			

measurement is classified using the stored templates. The classification outcome is compared to the classification of tracks in the vicinity of the new measurement to help with the association process. This approach makes use of the traditional automatic target recognition (ATR) process and thus has some of the same strengths and weaknesses. For instance, it can be sensitive to variations between testing data and training data. In other words, if a target encountered in the field does not match a stored template for that target, then this technique will perform poorly. The second technique does not rely on stored templates, but rather dynamically collects templates during the course of a mission. When a new measurement is collected, it is compared to the previously collected measurements that have been associated with tracks in the vicinity of the new measurement. The degree to which the new measurements match the previously collected measurements is used to aid the association of the new measurement to a track. This approach has the advantage of being able to eliminate mismatches between training and testing data, but it has a weakness in that it takes time to acquire enough measurements to become feasible. This is because the HRRGMTI profiles change rapidly with aspect angle. Optimum matching performance occurs when the difference in angle between two measurements is less than a degree. This would require over three hundred evenly spaced measurements to make this technique feasible at any aspect, even assuming a fixed elevation angle.

Each of the two techniques has advantages and disadvantages. A potential solution is to develop an algorithm that takes advantage of the strengths of each approach. For instance, the stored class templates could be used when no dynamic templates are available for a particular azimuth. When a dynamic template is available, it would be used instead. It will be key to determine a good point at which to switch between the two approaches.

In Phase I of this effort we have developed the class-oriented (traditional ATR) approach for signature-aided tracking using HRRGMTI profiles. We plan to extend our approach to dynamic templates in a Phase II effort. A block diagram of the class-oriented technique is shown in figure 1. The process begins with a batch of new measurements and a set of existing tracks. Validation regions about the tracks are set up, using kinematic information, to get a first filtering of the detections so that we do not consider associations that are physically impossible. For each detection in the validation region of a track, a kinematic
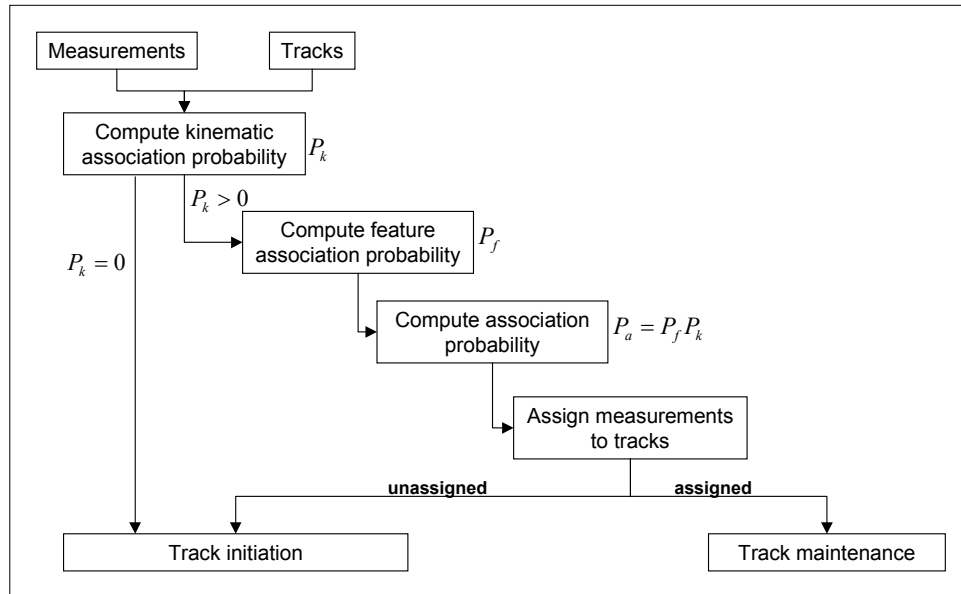

Figure 1: Signature-Aided Association

association probability and a signature association probability are calculated. In this discussion, we focus on the estimation of the signature association probability. To estimate the signature association probability, we first compute the probability that a vehicle in track belongs to a class (P_i) and the probability that a vehicle producing a new measurement belongs to the same class (P_{zi}). The track probabilities are recursively updated using each measurement that is associated with the track. The measurement probabilities are computed once for each measurement. The probability that a vehicle in track is from the same class that produced a given measurement (signature association probability) is simply given by

$$P_f = \sum_{i=1}^{N} P_i P_{zi}$$

where N is the number of classes. This probability can be combined with an estimate of the kinematic association probability (P_k) by simply multiplying the two numbers. As shown in figure 1, the resulting product can be used by the tracker to determine which detections should be associated with which tracks.

The assignment of detections to tracks is done by using a two-dimensional assignment matrix. There is one row in the matrix for each track and one column in the matrix for each detection. Additionally, a null row and column are considered for tracks that get no detections and detections that are assigned no tracks. The elements of the matrix consist of the association probabilities that are computed using the product of the kinematic and feature association probabilities. For the results presented in this document, assignment was performed using the Munkres Algorithm. This algorithm selects assignments so that the sum of the squares of the values in the matrix that correspond to the assignments is minimized. We have rerun these cases using a modified auction algorithm to perform the assignments and this produced nearly identical results.

2.1.1 Computation of P_{zi}

The probability that a measurement belongs to a given class (P_{zi}) would best be computed by using probability density functions that represent the likelihood of getting a particular signature given the target class. Due to the volatility of the signatures, this would need to be done in several angle bins. The likelihood that a given measurement belongs to that class could be determined by computing the value of the pdf at the point specified by the measurement. Unfortunately, since range profiles often have over 100 elements, such a pdf would be very difficult to construct due to a lack of data and difficulty representing a 100-dimensional, or larger, function. As an alternative, the actual density function can be approximated by assuming a general form that requires the specification of a small number of parameters. This can still be difficult. For instance, a Gaussian approximation to a 100-dimensional function would require the specification of 10,000 parameters for the covariance matrix and 100 parameters for the means.

We have approximated the multi-dimensional density function by creating a one-dimensional density function that represents the distribution of match scores when matching observations to a template range profile. The template range profile is created by taking the mean value of a set of range profiles in an angle bin. The one-dimensional density function is created by matching another set of range profiles against this template. The distribution of match scores provides a likelihood of getting a particular match given that the measurement belongs to a particular class. A different match-score pdf is created for each angle bin and each target class. In the next subsection, the matching algorithm that we used is discussed. We obtain a likelihood that a new measurement belongs to a particular class by computing the value of the matching density function at the point given by the match score which was found by matching the measurement to the template.

In phase II, a better approximation to the pdf will be created. Possible approaches include the creation of a pdf that assumes independent Gaussian distributions in each range bin[3] and the fitting of a chi-square distribution to the data[4].

The probability that the measured vehicle belongs to a given class can be computed using the likelihood obtained from the match-score pdfs:

$$P_{zi} = \frac{p_i(M_i)}{\sum_{j=1}^{N} p_j(M_j)}$$

where M_i is the match score when comparing a measurement to a template from a given class and $p_i(M_i)$ is the value of the pdf for class i at the match-score value.

In the case of dynamic templates, the match-score pdf described above cannot be built using only a few measurements, so we must assume a pdf. However, the assumed pdf can still be based on experimental data. For instance, an experiment can be performed in which measurements from a given vehicle are matched against templates from the same vehicle at nearly the same aspect angle. We will create such pdfs in Phase II using collected data. Previous analyses at Toyon have indicated that a Gaussian density function approximates the shape of these pdfs fairly well. Additionally, we have seen that many of these pdfs have a similar mean and variance. It would thus be fairly accurate to assume that the match-score pdfs in the dynamic template algorithm would be the same as the experimental pdfs.

2.1.2 Computation of P_i

The computation of the probability that a vehicle in track belongs to a given class based on all of the measurements collected thus far is computed using a recursive application of Bayes' Rule:

$$P_n(\omega_i) = \frac{P_{n-1}(\omega_i)p(s \mid \omega_i)}{\sum_{k=1}^{N} P_{n-1}(\omega_k)p(s \mid \omega_k)}$$

where ω_i is class i, N is the number of classes, s is the match score, $P_n(\omega_i)$ is the probability that the track belongs to class ω_i after n updates, and $p(s|\omega_i)$ is the likelihood of getting match score s given that the track belongs to class ω_i. We computed the values for $p(s|\omega_i)$ in this study using the match-score pdfs described in the previous subsection. In phase II, we will consider other approaches to computing $p(s|\omega_i)$.

2.1.3 Computation of P_k

We have employed a six-state Kalman filter for estimating the state of tracks. We do not use road-clamping or interacting multiple models. This is a fairly simple estimation procedure, but one that has proven to work well in a sparse target environment. In order to estimate P_k, we use the location of the detection and the estimated location of a track provided by the Kalman filter at the time of the detection. Additionally, the estimated errors provided by the track and sensor covariance matrices are used. This is done by creating a residual covariance matrix that represents the density function of the error in the estimate and the measurement. The likelihood of getting a particular detection is computed by finding the value of the residual density function at the location of the detection.

2.2 Analysis of range profile match surfaces

A key driver of performance for signature-aided tracking is the ability to determine the azimuth at which a target is viewed. The reason that this is the case can be understood by analyzing what happens when a measured profile is matched against templates at a variety of different angles. In order to do this, we must first create templates and observations and then select a matching procedure. Our templates were created by averaging sets of range profiles over a one-degree azimuth angle span. The range profiles used for the templates were extracted from SAR images collected in the MSTAR data collection. For the match surfaces shown here, we use this template data for our "observation" data as well. (However, in the simulations described below, the observation data was extracted from different MSTAR SAR images than those used to create the templates.)

The matching procedure that we selected was developed at AFRL and is used on the SHARP program. The matching algorithm used in SHARP was created by researchers at AFRL and is based on a linear regression fit of the input template and observation profile. Figure 2 illustrates the scoring method. Let $\mathbf{X}$ and $\mathbf{Y}$ represent vectors of template and observation RCS values, respectively. The vectors are ordered by range bin. The individual elements of $\mathbf{X}$ and $\mathbf{Y}$ are then regarded as coordinate pairs (i.e., the first coordinate pair is (first element of $\mathbf{X}$, first element of $\mathbf{Y}$) and so on), and a least squares line is fit to the coordinates. The vertical distance d from the observation profile RCS value to the least squares fit line is calculated for all RCS values in the observation profile vector. A score is then calculated from

$$Score = \sum_{i=1}^{N}(d_i)^2$$

where N is the number of range bins used in the scoring procedure. N was chosen as 70 range bins (+/- 23 feet from the profile centroid) to allow some signature overshoot even for larger targets. The template and observation are initially aligned by their respective centroids. The centroid is defined as the range bin at which the area underneath the range profile has been bisected. The template is then shifted +/- 5 range bins to allow for slight template/observation range misalignments, and a score is computed at each shift. The minimum score from each of these shifts is then regarded as the final discriminant score for that particular template/observation profile combination.

Given the above matching method, we can now compute what happens when we match templates at two different azimuths. Figure 3 contains a plot showing the match score obtained when T72 templates are matched against themselves at every azimuth. The resulting MSE scores give a two-dimensional match surface; the color coding indicates the height of the surface. In these match-surface plots, linear interpolation is used in regions of missing data so that the essential features of the figure are more evident.

Figure 3 shows that the best (lowest) scores are obtained when the azimuth at which the observation is collected is the same as the azimuth of the template. This is intuitively obvious and indicates that we will need to know the azimuth of each vehicle

Figure 2: AFRL's Linear Regression Matching Algorithm

Figure 3: Match Surface for T72 Matched Against T72

being tracked fairly accurately in order for FAT with HRR profiles to be most effective. Additionally, it indicates that the match scores themselves can be used to help compute the azimuth of the target. For instance, one could find the minimum value of the match surface within a window and use that as the selected azimuth. We discuss our procedure for estimating the target's azimuth in the next subsection.

Figure 3 also shows that there is a significant degree of symmetry in the T72. By symmetry we mean that profiles collected at a particular angle clockwise from the nose of the vehicle are similar to profiles collected at the same angle taken in a counter-clockwise direction from the nose of the vehicle. The dark blue line that goes from the upper left corner of the figure to the lower right corner indicates the presence of symmetry with respect to the centerline of the vehicle (a line going from the front of the vehicle to the back). Symmetry with respect to a line that is perpendicular to the centerline is also evident as seen by the shorter blue diagonals in the figure. Symmetry indicates that when using the best match score to compute azimuth, we could easily end up picking a symmetric angle rather than the true angle.

Symmetry may also play a role when developing the dynamic template approach. A possible replacement for a dynamically collected template at a particular azimuth may be a template created from profiles collected at the symmetric angle. We will explore this notion in a Phase II effort.

Figure 4 shows the match surface when the templates are M35 profiles matched against themselves. The M35 shows much less sensitivity to angle than the T72 in the forward half of the vehicle. This is indicated by the large rectangular regions of blue, and is due to the presence of a persistent scatterer that dominates the profile over a fairly large azimuth range. The M35 also shows much less symmetry than the T72.


Figure 4: Match Surface for M35 Matched Against M35

Figure 5 shows the match surface when M35 templates are matched against T72 templates. The results indicate that the observations do not match very well even when they are at the same azimuth. The results also indicate that when both templates are near broadside, the match is fairly good. This is because the profiles tend to turn into a single lump near broadside for all types of vehicles. In response to this, we have set up a broadside exclusion zone so that when we are near broadside, we rely on kinematic information only and do not include the signature information. Note that the radar is least likely to detect a target at broadside due to the very slow range rate relative to the clutter.

Figure 5: Match Surface for M35 Matched Against T72

2.3 Azimuth estimation procedure

To estimate the azimuth of a vehicle, we make use of road network information and the match surfaces described in the previous section. In previous analyses, we have used state estimates provided by the tracker instead of road network information, but this resulted in worse performance. In the future, we would like to combine both techniques due to the potential fragility of the road network approach. This fragility comes from the fact that there are errors in the road data and vehicles do not always travel on roads.

We begin our azimuth estimation by clamping the detection to the nearest road segment that is in the same range bin, as shown in figure 6. We can then compute the direction that the vehicle is traveling using the Doppler measurement and knowledge of the sensor and road network geometry. Using this information, and knowledge of the radar's position, we can estimate target azimuth. If the nearest road segment is the correct road segment, this gives us a good estimate of azimuth. However, if the vehicle is near an intersection, the correct road segment may not be the nearest segment. This is because there are typically large azimuth measurement errors present that are larger than the distance between roads when the target is near an intersection. We thus need a procedure for recognizing that we have a bad azimuth estimate and a method for dealing with the situation. We have chosen to recognize the problem by storing the best match that is obtained when matching an observation to all of the template classes. If the best match is worse than a threshold


Figure 6: Geometry of Road-Clamping Algorithm

value, the azimuth is declared invalid and we search over a much larger window to find the correct azimuth based solely on the profile match.

The search for the correct azimuth is done by comparing the observation to all of the templates from each class over a large azimuth span. The azimuth of the template that produced the best match score is selected as the correct azimuth. This is a simple approach that we plan to improve upon in the future. For instance, a joint pdf that captures the likelihoods of both the kinematic and feature matches could be created and used to compute an estimate of target azimuth.

2.4 Testbed development

Quantification of the performance of the signature-aided-tracking approach described above must account for the dynamic nature of the problem. Quantifying performance in a meaningful manner requires the use of a computer simulation. We have modified Toyon's SLAMEM™ simulation in order to use it as a testbed for our algorithm development. The key pieces of this development are shown in the block diagram of figure 7.
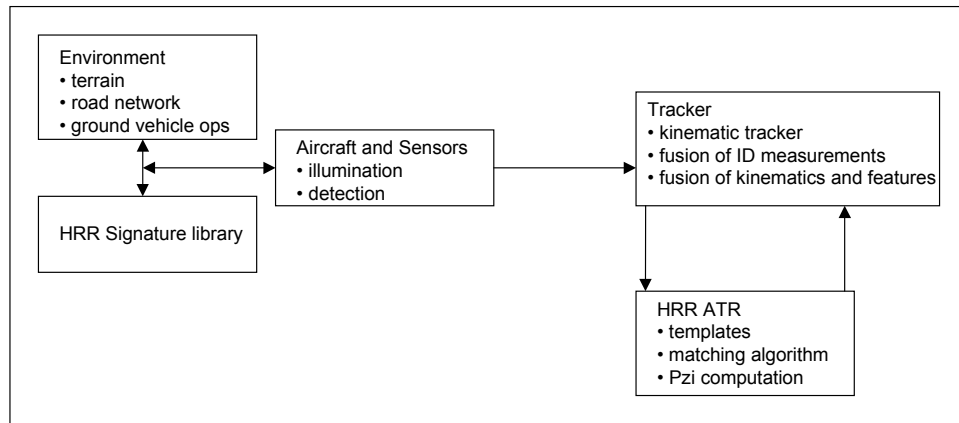

Figure 7: Analysis Testbed

The SLAMEM™ simulation contains a model of ground vehicle operations, road networks, and terrain. In this effort we added the box labeled "HRR signature library." When a vehicle is detected at a particular azimuth and elevation, a profile is selected from the signature library that is closest to the observed geometry. This signature is then attached to the detection and passed along to the tracker. This library was created using MSTAR data. In a Phase II effort we expect to make use of a developing signature library that is being created at AFRL/SN.

A model of the aircraft and its sensor are used to determine when a set of ground vehicles are illuminated. For those vehicles that are illuminated, a clear-line-of-sight check is done to determine if the vehicle is masked by terrain. If the vehicle is not masked by terrain, its speed projected along the line-of-sight is computed to determine if it is above the minimum detectable velocity of the sensor. If it exceeds the MDV, a random number is drawn and compared to a probability of detection to determine if the vehicle is detected. If the vehicle is detected, its state is mapped into range, azimuth, and Doppler. Appropriate errors are added to each measurement component to reflect measurement errors. If more than one vehicle is in the same range, azimuth, and Doppler resolution bin, only one detection is created. The corrupted detections, along with the range profile, are sent to the tracker.

The tracking module performs the state estimation of tracks given a set of measurements, the association of the detections to tracks, and the accumulation of evidence from multiple ID measurements that have been collected over time. State estimation is done using a six-state Kalman Filter. The measurements are associated to tracks by forming an assignment matrix that is filled with association likelihoods as discussed in Section 2.1.1. The accumulation of evidence using multiple measurements is performed by applying Bayes' Rule.

The tracker interacts with an HRR ATR module that computes the class probabilities for each measurement (P_{zi}). The ATR has stored templates for each class, created using MSTAR data. The data used in the signature library to represent measurements was taken from different target chips than those that were used to create the templates in the ATR module. The

ATR module computes the probabilities as discussed in Section 2.1.1. When we use dynamic templates for SAT, this module will contain software that collects the measurements and creates the dynamic templates.

2.5 Evaluation of signature-aided tracking

In order to quantify the performance improvement gained when using SAT, we have analyzed track performance in the scenario that is depicted in figure 8. The figure shows the location of two aircraft that are located on orbits approximately 100 km from the center of the region. The aircraft are flying at an altitude of 66,000 ft and have two thirteen-minute scanning legs and two three-minute turns in their orbits. Two T72 tanks are located within the 12 km x 12 km box shown in the figure. The two aircraft have GMTI sensors which scan the region in an attempt to maintain knowledge of the location of each tank. Also traveling on the roads within the area are a number of M35 background vehicles. We consider performance for a range of the number of background vehicles between zero and one hundred.


Figure 8: Scenario Used for Analysis

The measure of performance that we will use to quantify the effectiveness of SAT is the duration of a track. The track duration is not simply the length of time that the track exists in the tracker. Rather, it is the length of time that we are able to maintain knowledge about the location of a particular vehicle. This measure is illustrated in figure 9. The figure shows a cartoon where the shaded squares represent vehicles at different times. The shaded bar at the bottom of the figure indicates which vehicle is the source of the detections used to update a track that starts with the blue vehicle. The track starts out gray, but when the black and gray vehicles come close together, some black detections are used to update the track. Eventually, the black vehicle separates from the gray vehicle and the track reverts back to pure gray detection updates. One metric that we consider for track duration is the length of time from track initiation to the time at which a single detection from another vehicle is used to update the track. This is noted as "perfect association" in figure 9. This is a harsh metric because a track may get updated once by a nearby vehicle, yet still be tracked for an extended period after the one false update. To account for this, we consider a window of time such that if the track is updated with a detection from the original vehicle during this time, then the track duration continues. This is illustrated by the continuation of the duration shown in the figure after the initial confusion with the black vehicle. Eventually, the white vehicle becomes confused with the gray vehicle and the track is transferred to it. At this point we consider the duration of the track to have ended.
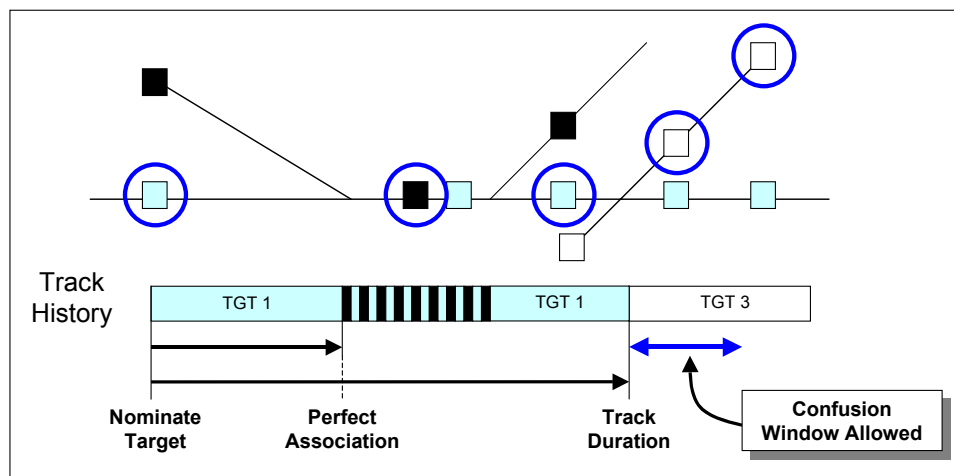

Figure 9: Measure of Performance

The results for the kinematic tracker when using the above measures of performance for a case with 50 background vehicles are plotted in figure 10. The figure shows the fraction of tracks that have the duration specified by the horizontal axis. The fraction of tracks was computed using multiple Monte Carlo trials of the simulation. A curve for both the perfect duration and windowed duration are shown. The results show that for this background density, the vehicles can only be tracked for a short

Figure 10: Kinematic Tracking Performance

period of time before a mistake is made. For example, a reasonable goal is to be able to track these vehicles for more than thirty minutes. Since we are far short of this goal, even for a small number of background vehicles, it is clear that assistance from signatures or something else is needed.

By using the signatures of the vehicles as described earlier, the track duration can be significantly increased. This is shown in figure 11 where the track duration of 50% of the tracks is plotted versus the number of background vehicles. The bottom curve indicates the performance of the tracker when using only kinematic information. The middle curve indicates the performance when employing SAT. Note that performance is significantly improved over the kinematic-only case, but still is not quite as good as we would like to see at the higher number of background vehicles. The top curve in the figure shows the performance when using SAT with perfect knowledge of vehicle azimuth. This is a hypothetical case that can only be analyzed given ground truth, but it serves to illustrate that one of the key drivers of performance is the ability to correctly estimate the azimuth of a vehicle that is detected.

Our results indicate that there is great potential for signature-aided tracking; however, there are many technical problems that must be addressed prior to the development of a deployable system. For instance, we have only considered two classes of vehicles in this example. In many situations many more classes will need to be considered in addition to an unknown class. Furthermore, we were able to get fairly good results with the SAT approach because we were able to use road information to help with the estimation of vehicle azimuth. Vehicles will not always travel on roads and thus an approach for better estimating azimuth without road information needs to be developed. Additionally, the vehicles that we were tracking did not stop at any time while we were tracking them. Since vehicles control where and when they stop, we need to develop logic that allows us to track them through move-stop transitions.

Figure 11: Signature-Aided-Tracking Performance

ACKNOWLEDGMENTS

REFERENCES

1. M. T. Fennell, R. P. Wishner, "Battlefield Awareness Via Synergistic SAR and MTI Exploitation", *IEEE AES Systems Magazine*, Vol. 13, No. 2, pp. 39-45, 1998.
2. K. J. Sullivan, M. R. Meloon, M. B. Ressler, "Estimate of Continuous Tracking Performance using SAR and MTI Surveillance," Toyon Research Corporation (T245C7001), 15 October 1998.
3. R. A. Mitchell, J. J. Westerkamp, "Robust Statistical Feature Based Aircraft Identification", *IEEE AES Systems*, Vol. 35, No. 3, July 1999.
4. D. Miklovic, C. Holt, S. Schmidt, "Low Variability Features for Signature Aided Tracking", 55th ATRWG, AETC.